

Strongly Unforgeable Signatures and Hierarchical Identity-based Signatures from Lattices without Random Oracles

Markus Rückert*
markus.rueckert@cased.de

Cryptography and Computeralgebra
Department of Computer Science
TU Darmstadt

Abstract. We propose a variant of the “bonsai tree” signature scheme, a lattice-based *existentially unforgeable* signature scheme in the standard model. Our construction offers the same efficiency as the “bonsai tree” scheme but supports the stronger notion of *strong unforgeability*. Strong unforgeability demands that the adversary is unable to produce a new message-signature pair $(\mathbf{m}, \mathfrak{s})$, even if he or she is allowed to see a different signature \mathfrak{s}' for \mathbf{m} .

In particular, we provide the first treeless signature scheme that supports strong unforgeability for the post-quantum era in the standard model. Moreover, we show how to *directly* implement identity-based, and even hierarchical identity-based, signatures (IBS) in the same strong security model without random oracles. An additional advantage of this direct approach over the usual generic conversion of hierarchical identity-based encryption to IBS is that we can exploit the efficiency of ideal lattices without significantly harming security.

We equip all constructions with strong security proofs based on mild worst-case assumptions on lattices and we also propose concrete security parameters.

Keywords Post-quantum cryptography, lattice cryptography, digital signatures, identity-based cryptography, standard model

* This work was supported by CASED (www.cased.de).

1 Introduction

Digital signature schemes are the cornerstone of e-business, e-government, software security, and many more applications. Their importance is likely to grow in the future as more and more everyday tasks and processes are computerized. With identity-based signature schemes (IBS), motivated by Shamir [34], one can get rid of public-key infrastructures. The public key is replaced with a unique identifier string, such as an e-mail address, and the secret key is “extracted” by a trusted party for this identifier. In hierarchical identity-based signatures (HIBS), this concept is generalized so that each party can act as a key extraction authority for its subordinates.

There are two classes of signature schemes. The first comprises *tree-based* and stateful Merkle signature schemes [29] with a limited signature capacity. Such schemes can be solely based on the security of hash functions. The drawback is that they require an inefficient key generation phase, where all signatures need to be prepared in advance. Furthermore, its statefulness poses a synchronization problem as soon as more than one computer, or process thread, is supposed to issue signatures with the same secret key.

The second class contains *treeless* constructions that are typically more efficient and allow for an unlimited number of signatures without a complex setup phase. Currently, we mainly use schemes that fall into the second category because they are easier to handle. Most of them rely on the hardness of factoring or computing discrete logarithms.

Alternatives for the post-quantum era can be based on the hardness of the decoding problem in error correcting codes, on the hardness of solving nonlinear multivariate equation systems, or on the hardness of lattice problems. Refer to [7] for an overview of each field. Basically, all three alternatives rely on the hardness of certain *average-case* problems and, at first, it is unclear how to generate hard instances of these problems. More precisely, we always need to know a “hard” distribution of keys that admits efficient key generation. Unlike with multivariate or code-based cryptography, lattice-based constructions have a “trust anchor” in the form of Ajtai’s worst-case to average-case reduction [2] that is not found anywhere else in cryptography. It states that solving a certain average-case problem, which is relevant in cryptography, implies a solution to a related worst-case problem. Although this may sound purely theoretical, it is of great practical value as keys that are chosen uniformly at random already provide *worst-case* security guarantees. The hardness of this underlying worst-case problem is also plausible as the best known algorithm to solve the relevant lattice problems requires exponential time [3].

Another advantage of lattice-based cryptography over the alternatives is that there is a whole range of provably secure signature schemes. In the random oracle model there are schemes due to Gentry, Peikert, and Vaikuntanathan [16]; Stehlé, Steinfeld, Tanaka, and Xagawa [35]; and Lyubashevsky [26]. As for the standard model, there are the works of Lyubashevsky and Micciancio [27] (tree-based); and Cash, Hofheinz, Kiltz, and Peikert [12].

However, there is a gap in this range because none of the above schemes is *stateless*, provably secure in the *standard model*, and *strongly unforgeable*. Strong unforgeability under chosen message attacks (SU-CMA) is stronger than existential unforgeability (EU-CMA) in the sense that the adversary is not artificially restricted by the security model. In EU-CMA, the adversary is forced to output a signature for a *fresh message* \mathbf{m}^* after seeing signature for messages $\mathbf{m}_i \neq \mathbf{m}^*$ of his or her choice. The SU-CMA adversary is also allowed to output a *fresh signature* for one of the \mathbf{m}_i .

Strong unforgeability is interesting in both, theory and practice. Consider generic transformations from CPA to CCA2 security in the standard model, e.g., Dolev, Dwork, and Naor [13] or Boneh, Canetti, Halevi, and Katz [9]. They typically involve a strongly unforgeable signature scheme to make the ciphertext authentic and non-malleable. An EU-CMA signature of the CPA ciphertext may already provide some security against CCA1 adversaries but a CCA2 attack would certainly still succeed. Another reason is the construction of ID-based blind signatures due to Galindo, Herranz, and Kiltz [14].

As a practical example, consider an access control protocol where you may delegate certain rights to another party by signing a description for these rights with a signature \mathfrak{s} . You want to be able to revoke them at any time in the future via an online revocation system. The rights are revoked as soon as the online system has \mathfrak{s} in its database. If the signature scheme is only EU-CMA secure, the delegator can construct another signature \mathfrak{s}^* for the same set of rights and present this token instead of \mathfrak{s} — the revocation mechanism breaks down.

Notice that there are generic transformations from EU-CMA to SU-CMA. They typically only apply to a certain small subclass of signature schemes, e.g., Boneh-Shen-Waters [10]. Recently, Bellare and Shoup [6] proposed an unrestricted transformation. However, all lose efficiency compared to the underlying EU-CMA scheme because they require multiple signing steps.

Similarly, HIBS can be constructed from any signature scheme, e.g., via the certification approach [17] by Gentry and Silverberg or via Append-Only Signatures due to Kiltz, Mityagin, Panjwani, and Raghavan [20]. Here, the key extraction authority basically emulates a public-key infrastructure: It generates a new public key for each identity and issues a certificate, binding the key to the identity. In consequence, the resulting identity-based signatures grow by the size of one public key, which will not be practical in our case.

Our Contribution. Table 1 compares our result with the current state-of-the-art for lattice signatures, including the typical improvements with ideal lattices [32,30]. All previously known SU-CMA schemes are either stateful, causing, e.g., synchronization problems, or they require random oracles. Using random oracles is discouraged by the works of Canetti, Goldreich, and Halevi [11], as well as by the more practical work of Leurent and Nguyen [24]. The only known scheme that directly provides SU-CMA security in the standard model is stateful and has

a large secret key.¹ Our construction in Section 4.1 offers the same complexity as [12]. In our case, signing involves a simple additional linear algebra step that can be pre-computed. Thus, we achieve a stronger security notion without additional cost.

Scheme	Stateless	Standard model	SU-CMA	Public key	Secret Key	Signature
[16] with [35]	Yes	No	Yes	$\tilde{\mathcal{O}}(n)$	$\tilde{\mathcal{O}}(n^2)$	$\tilde{\mathcal{O}}(n)$
[26]	Yes	No	Yes	$\tilde{\mathcal{O}}(n)$	$\mathcal{O}(n)$	$\tilde{\mathcal{O}}(n)$
[12]	Yes	Yes	No	$\tilde{\mathcal{O}}(n)$	$\tilde{\mathcal{O}}(n^2)$	$\tilde{\mathcal{O}}(n)$
[27] with [29]	No	Yes	Yes	$\mathcal{O}(n)$	$n^{\mathcal{O}(1)}$	$\tilde{\mathcal{O}}(n)$
Section 4.1	Yes	Yes	Yes	$\tilde{\mathcal{O}}(n)$	$\tilde{\mathcal{O}}(n^2)$	$\tilde{\mathcal{O}}(n)$

Table 1. Comparison of the properties of current lattice-based signature schemes.

The situation is quite similar for HIBS. With the exception of Libert and Quisquater [25], previous results deal only with existentially unforgeable HIBS. They typically provide hierarchical identity-based encryption (HIBE) and then apply a generic conversion, e.g., [21], to obtain an HIBS. Both, HIBE and HIBS, can be classified as selective-ID or adaptive-ID secure. Selective-ID security forces the adversary to name its target identity before seeing the public key. In the adaptive case, it may output a forgery for any identity. So far, lattice-based IBE schemes either support selective-ID security [12,1] in the standard model or they support adaptive-ID security [16] in the random oracle model. In contrast to our constructions, using the HIBE to HIBS conversion with the above HIBE schemes requires a subexponential time quantum reduction in the security proof when using efficient ideal lattices [35].

Note that using the generic construction [17] in conjunction with our SU-CMA signature scheme (Section 4) would work and the result would be a strongly-unforgeable HIBS. However, the resulting signature size would be completely impractical.

Hence, in addition to the first stateless standard-model SU-CMA signature scheme from lattices in Section 4.1, we provide the first direct lattice-based constructions for adaptive-ID secure and strongly unforgeably HIBS in the random oracle model in Section 3.2 and for strongly unforgeable HIBS in the standard model in Section 4.2. Our constructions in Section 4 rely on Chameleon hash functions [23]. The security proofs involve a generic transformation to SU-CMA from a slightly weaker notion, which was not explicitly known before (cf. Section 2). Due to space restrictions, some of the proofs had to be moved to the full version [33].

Our Modifications to [12]. For those familiar with [12], we give a brief overview of the changes that are necessary to achieve SU-CMA security. In [12], signatures

¹ Trade-offs are possible but the key generation complexity is $n^{\mathcal{O}(1)}$, typically for a large polynomial.

are short vectors \mathfrak{s} that satisfy $\mathbf{A}\mathfrak{s} \equiv \mathbf{0} \pmod{q}$, i.e., they are in the q -ary lattice $\Lambda_q^\perp(\mathbf{A})$. An adversary may succeed in breaking SU-CMA security of this scheme by simply asking for a signature \mathfrak{s} for a message \mathbf{m} and then return $(\mathbf{m}, -\mathfrak{s})$ as its forgery. Such an answer is useless in the reduction because the simulator had to know a trapdoor for $\Lambda_q^\perp(\mathbf{A})$ beforehand.

Instead, we let the signature algorithm sample short vectors from a random coset $\{\mathbf{x} : \mathbf{A}\mathbf{x} \equiv \mathbf{y} \pmod{q}\}$. The approach is similar to [16] but with a fixed \mathbf{y} that is part of the public key. In the simulation, we can prepare \mathbf{y} such that we know a corresponding signature \mathfrak{s} that is used to simulate the signature oracle. Now, the adversary against SU-CMA security needs to output a different short vector \mathfrak{s}^* from the same coset. This, however enables the simulation to find a short vector $\mathfrak{s} - \mathfrak{s}^*$ in $\Lambda_q^\perp(\mathbf{A})$ and solve the underlying problem.

2 Preliminaries

The security parameter is n . The statement $x \leftarrow_{\mathfrak{s}} X$ means x is chosen uniformly at random from X . With $x \sim \Delta(X)$, we denote that x is chosen according to a distribution Δ over X . The concatenation of strings, vectors, and matrix columns is done via \circ . Furthermore, $x \sqsubset y$ means x is a prefix of y and \emptyset is the empty string. Lower-case boldface identifies vectors and upper-case boldface denotes a matrix. For a given matrix $\mathbf{X} \in \mathbb{R}^{m \times m}$, we write $\tilde{\mathbf{X}}$ for its Gram-Schmidt orthogonalization.

2.1 Security Models

Throughout the paper we stick to the following notation. The length of identities is κ , the message length is λ , and ℓ is the hierarchy depth, meaning that there are $\ell + 1$ levels in the hierarchy tree. With $\{x_i\}_1^m$ we denote the set $\{x_1, \dots, x_m\}$. The subsequent paragraphs deal with the specification of strongly unforgeable signature schemes $\text{DSig} = (\text{Kg}, \text{Sign}, \text{Vf})$ and hierarchical identity-based signature schemes $\text{HIBS} = (\text{Kg}, \text{Extract}, \text{Sign}, \text{Vf})$.

Signature Schemes. We follow the standard specification for digital signature schemes: $\text{Kg}(1^n)$ outputs a private signing key sk and a public verification key pk ; $\text{Sign}(\text{sk}, \mathbf{m})$ outputs a signature \mathfrak{s} under sk for the message \mathbf{m} ; $\text{Vf}(\text{pk}, \mathfrak{s}, \mathbf{m})$ outputs 1 iff \mathfrak{s} is a valid signature on \mathbf{m} under pk .

Most schemes are proven to be existentially unforgeable under chosen message attacks (EU-CMA), but we will consider the stronger notion of strongly unforgeability under chosen message attacks (SU-CMA) as described in the following experiment, where \mathcal{H} is a family of random oracles.

Experiment $\text{Exp}_{\mathcal{A}, \text{DSig}}^{\text{SU-CMA}}(n)$

$\mathcal{H} \leftarrow \mathcal{H}(1^n)$; $(\text{sk}, \text{pk}) \leftarrow \text{Kg}(1^n)$

$(\mathbf{m}^*, \mathfrak{s}^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot), \mathcal{H}(\cdot)}(\text{pk})$

Let $(\mathbf{m}_i, \mathfrak{s}_i)$ be the answer returned by $\text{Sign}(\text{sk}, \cdot)$ on input \mathbf{m}_i , for $i = 1, \dots, k$.

Return 1 iff $\text{Vf}(\text{pk}, \mathbf{m}^*, \mathfrak{s}^*) = 1$ and $(\mathbf{m}^*, \mathfrak{s}^*) \notin \{(\mathbf{m}_1, \mathfrak{s}_1), \dots, (\mathbf{m}_k, \mathfrak{s}_k)\}$.

DSig is (t, q_S, q_H, ϵ) -strongly unforgeable if there is no t -time adversary that succeeds with probability $\geq \epsilon$ after making $\leq q_S$ signature oracle queries and $\leq q_H$ random oracle queries. In the standard model, we leave out H and q_H .

The difference to the EU-CMA model is that the adversary in the SU-CMA model even wins if it outputs a new signature for a message that it already knows a signature for. In the EU-CMA model, the adversary is forced to output a forgery for a “fresh” message. Instead of directly providing SU-CMA security in our main constructions in Section 4, we use the weaker notion of strong unforgeability against static message attacks (SU-SMA). Here, the adversary submits all messages $\mathbf{m}_1, \dots, \mathbf{m}_{q_S}$ before seeing the public key and the corresponding signatures. Then, we use a generic transformation to achieve full security.

HIBS Schemes. The specification for HIBS schemes is a straightforward generalization of that for digital signature schemes. The main difference is that there are no per-signer verification keys but rather a shared verification key for the entire system. Moreover, the signer’s public key is easily computable from a unique user identification string ID over a binary alphabet. The corresponding secret signing key is “extracted” by a trusted authority using a master secret key. The hierarchy is modeled by letting identities be a concatenation of per-level identifiers with decreasing rank, i.e., $ID = ID_0 \circ ID_1 \circ ID_2$ describes an identity on level 2 with parent identity $ID_0 \circ ID_1$, whose parent identity is the master identity ID_0 .

More formally: $\text{Kg}(1^n)$ outputs a master private key msk and a master public key mpk . The master identity is the empty string \emptyset ; $\text{Extract}(\text{sk}_{ID^*}, ID)$ outputs a secret signing key sk for ID if $ID^* \sqsubset ID$, otherwise \perp ; $\text{Sign}(\text{sk}_{ID}, ID, \mathbf{m})$ outputs a signature \mathfrak{s} under sk_{ID} for \mathbf{m} ; $\text{Vf}(\text{mpk}, ID, \mathfrak{s}, \mathbf{m})$ outputs 1 iff \mathfrak{s} is a valid signature on \mathbf{m} for the given identity and master public key.

The security models for HIBS and ordinary signatures are tightly related with the exception that one has to deal with the additional key extraction mechanism. We consider two variants, selective-ID security (similar to selective-secure HIBE [8]) and the stronger notion of adaptive-ID security. In both models, the adversary can query a key extraction oracle E , a signature oracle, and an optional random oracle. The experiment $\text{Exp}_{\mathcal{A}, \text{HIBS}}^{\text{SU-CMA-SelectiveID}}$ describes selective-ID security, where the adversary has to fix an identity ID^* before seeing the master public key. He is then forced to output a forgery for ID^* . The adversary gets secret keys for all identities that are not a prefix of ID^* . Furthermore, it can query a signature oracle $S(ID, \mathbf{m})$ with arbitrary identities and messages.

Experiment $\text{Exp}_{\mathcal{A}, \text{HIBS}}^{\text{SU-CMA-SelectiveID}}(n)$
 $H \leftarrow \mathcal{H}(1^n)$; $(ID^*, \text{state}) \leftarrow \mathcal{A}(1^n)$; $(\text{msk}, \text{mpk}) \leftarrow \text{Kg}(1^n)$
 $(\mathbf{m}^*, \mathfrak{s}^*) \leftarrow \mathcal{A}^{E(\text{msk}, \cdot) \setminus \{\sqsubset ID^*\}, S(\cdot, \cdot), H(\cdot)}(\text{mpk}, \text{state})$
Let $\{(ID_i, \mathbf{m}_i, \mathfrak{s}_i)\}_1^k$ be the query-answer tuples for S .
Return 1 iff $\text{Vf}(\text{mpk}, ID^*, \mathfrak{s}^*, \mathbf{m}^*) = 1$ and $(ID^*, \mathbf{m}^*, \mathfrak{s}^*) \notin \{(ID_i, \mathbf{m}_i, \mathfrak{s}_i)\}_1^k$.

In the stronger model of adaptive-ID security, the adversary can output a forgery for any identity, for which he has never queried a prefix to the extraction oracle before.

Experiment $\text{Exp}_{\mathcal{A}, \text{HIBS}}^{\text{SU-CMA-AdaptiveID}}(n)$
 $\text{H} \leftarrow \mathcal{H}(1^n)$; $(\text{msk}, \text{mpk}) \leftarrow \text{Kg}(1^n)$
 $(\text{ID}^*, \text{m}^*, \text{s}^*) \leftarrow \mathcal{A}^{\text{E}(\text{msk}, \cdot), \text{S}(\cdot, \cdot), \text{H}(\cdot)}(\text{mpk})$
Let $\{(\text{ID}_i, \text{m}_i, \text{s}_i)\}_1^k$ be the query-answer tuples of S ;
Let $\{\text{ID}_j\}_1^l$ be the query-answer pairs of E .
Return 1 iff $\forall (\text{mpk}, \text{ID}^*, \text{s}^*, \text{m}^*) = 1$
and $(\text{ID}^*, \text{m}^*, \text{s}^*) \notin \{(\text{ID}_i, \text{m}_i, \text{s}_i)\}_1^k$
and $\{\text{ID}_j\}_1^l \ni \text{ID}_j \not\subseteq \text{ID}^*$.

HIBS is $(t, q_E, q_S, q_H, \epsilon)$ -strongly unforgeable under chosen message and selective (adaptive) identity attacks if there is no t -time adversary that succeeds with probability $\geq \epsilon$ after making $\leq q_E$ extraction queries, $\leq q_S$ signature oracle queries, and $\leq q_H$ random oracle queries in the respective experiment. Again, we leave out H and q_H in the standard model.

In Section 4, we will provide instantiations secure against static message attacks (SMA) and then use the following transformation to achieve CMA security.

From SMA to CMA. Krawczyk and Rabin [23] proposed Chameleon hashes to be hash functions with a trapdoor and the following properties. 1) The function $\text{C} : D \times E \rightarrow R$ is chosen from a family \mathcal{C} of Chameleon hashes along with a secret trapdoor t . 2) In order to sample from the distribution $(d, e, \text{C}(d)) \in D \times E \times R$, we can do one of two things. Either we run C on the given document d and a randomness $e \sim \Delta(E)$ (efficiently sampleable), or we apply an inversion algorithm $e \leftarrow \text{C}_t^{-1}(r, d)$ on a given image $r \in R$ and a target document $d \in D$. Thus, we obtain a randomness e such that $\text{C}(d) = (e, r)$. It is important that the resulting distributions are within negligible statistical distance. Note that whenever we need the Chameleon hash to map to a certain range $\neq R$, we can compose it with an arbitrary collision resistant hash function. As for their realization, Krawczyk and Rabin claim in [22] that Chameleon hash functions exist if there are claw-free trapdoor permutations. Interestingly, they can be easily implemented with the lattice-based trapdoor function in [16] as observed in [12].

A helpful fact about Chameleon hash functions is that if they exist, then there is a generic transformation from EU-SMA to EU-CMA signatures. This was known since [23] and it is proven in [19]. We show that the transformation also works for SU-SMA to SU-CMA in the full version [33]. Observe that it is also applicable to selective-ID secure HIBS schemes as it only affects the way the signature oracle is simulated for the challenge identity.

Lemma 1. *SU-SMA implies SU-CMA if Chameleon hash functions exist.*

2.2 Lattices

In this work, we deal only with full-rank q -ary lattices, i.e., lattices that represent the kernel of the linear map $\mathbf{x} \mapsto \mathbf{A}\mathbf{x} \bmod q$ for a prime modulus q and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. These lattices are denoted with $\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} \equiv \mathbf{0}\}$

(mod q)). As with all lattices of dimension $m \geq 2$, they have infinitely many bases. A basis of $\Lambda_q^\perp(\mathbf{A})$ is a matrix $\mathbf{B} \in \mathbb{Z}^{m \times m}$, such that $\Lambda(\mathbf{B}) := \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^m\}$ is equal to $\Lambda_q^\perp(\mathbf{A})$. The quantity $\det(\Lambda) = |\det(\mathbf{B})|$ (for any basis) is a lattice constant. The main computational problem in q -ary lattices is the “short integer solution” problem SIS. It is parameterized with positive integers $n, m = \text{poly}(n), q = \text{poly}(n)$, and a real norm bound ν and it is formulated as an average-case problem: Given a uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, find a non-zero $\mathbf{v} \in \Lambda_q^\perp(\mathbf{A})$ with $\|\mathbf{v}\|_2 \leq \nu$. Ajtai showed in [2] that this problem is at least as hard as finding short vectors in *all* lattices of dimension n , i.e., solving a related worst-case problem. A recent improvement to this reduction can be formulated as follows.

Theorem 1 (Worst-case to Average-case Reduction [16] (informal)). *If there is a polynomial time algorithm that breaks $\text{SIS}(n, m, q, \nu)$ for $q \geq \nu \omega(\sqrt{n \log(n)})$, $\nu = \text{poly}(n)$ with non-negligible probability, then there is a polynomial time algorithm that finds short non-zero vectors, which are only a $\gamma \geq \nu \tilde{O}(\sqrt{n})$ factor longer than the shortest vector, in all lattices of dimension n .*

In cryptography, we typically hand over \mathbf{A} , or a “bad” basis with long vectors, as the public key and keep a “good” (short) basis as our secret. The length of a basis is $\|\mathbf{B}\| := \max_{i=1, \dots, m} \|\mathbf{b}_i\|_2$. This principle goes back to Ajtai. The most recent improvement for generating such a matrix \mathbf{A} together with a particularly short trapdoor \mathbf{T} for SIS is due to Alwen and Peikert [5].

2.3 Bonsai Trees

The notion of “bonsai trees” on lattices is introduced in [12] in analogy to arboriculture. An arborist always starts with a certain amount of *undirected*, i.e., random, natural growth that he cannot control. Then, he applies his tools and starts cultivating individual branches to achieve the desired looks via *directed* growth. The arborist is successful if the resulting tree still looks sufficiently natural to the observer. Once cultivated, a branch can easily be *extended* to form more directed growth without too much additional care. Instead of extending directed growth, the arborist can also generate a *randomized* offsprings, which can be given to another arborist that can easily cultivate them by *extending* growth. The offsprings hide the first arborist’s work and the employed techniques. We formalize these concepts in the context of lattices. A (binary) bonsai tree is generated out of a root \mathbf{A}^* and branches $\mathbf{A}_i^{(b)} \in \mathbb{Z}_q^{n \times m_i}$, $b \in \{0, 1\}, i \leq k \leq \text{poly}(n)$, that are statistically close to uniform. The entire tree is the set $\{\mathbf{A}^* \circ \mathbf{A}_1^{(x_1)} \circ \dots \circ \mathbf{A}_k^{(x_k)} : \mathbf{x} \in \{0, 1\}^{\leq k}\}$.

Proposition 1 (Directed Growth). *Let $\delta > 0$ be any fixed real constant and let $q \geq 3$ be odd. There is a polynomial time algorithm $\text{ExtLattice}(\mathbf{A}_1, m_2)$ that, given uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m_1}$ for any $m_1 \geq (1 + \delta)n \log_2(q)$ and $\text{poly}(n)$ -bounded $m_2 \geq (4 + 2\delta)n \log_2(q)$, outputs $(\mathbf{A}_2 \in \mathbb{Z}_q^{n \times m_2}, \mathbf{S} \in \mathbb{Z}^{m \times m})$, where $m = m_1 + m_2$, such that $\mathbf{A} = \mathbf{A}_1 \circ \mathbf{A}_2$ is within negligible statistical distance*

of uniform; \mathbf{S} is a basis of $\Lambda_q^\perp(\mathbf{A}_1 \circ \mathbf{A}_2)$; $\|\mathbf{S}\| \leq L = Cn \log_2(q)$ with overwhelming probability; and $\|\tilde{\mathbf{S}}\| \leq \tilde{L} = 1 + C\sqrt{(1+\delta)n \log_2(n)} \leq 1 + C\sqrt{m_1}$ with overwhelming probability.

The proposition reflects the most recent result on lattice trapdoors from [4]. In the following, we will use $C = 20$, $\delta = 1$ for simplicity and assume that \mathbf{A}_2 is uniformly random instead of within negligible statistical distance from uniform. The resulting key sizes can be optimized by taking δ close to 0 instead. This results in a less uniform distribution of \mathbf{A}_2 but it is still within distance $m_2 q^{-\delta n/2}$ from uniform. The interpretation in terms of arboriculture is generating “directed growth” out of “undirected growth” because one starts with some random growth \mathbf{A}_1 and cultivates a branch $\mathbf{A}_1 \circ \mathbf{A}_2$ along with a trapdoor \mathbf{T} , which is the arborist’s journal or a trace of his work. However, the observer cannot distinguish undirected growth from directed growth.

An important observation is that knowing a trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ implies knowing a trapdoor for all $\mathbf{A}' \in \mathbb{Z}_q^{n \times m'}$, $m' \geq m$, when $\mathbf{A} \sqsubset \mathbf{A}'$. This is because one can apply the trapdoor in dimension m and then pad the resulting vector with zeros to solve SIS in dimension m' .

Proposition 2 (Extending Control). *There is polynomial time algorithm $\text{ExtBasis}(\mathbf{S}_1, \mathbf{A} = \mathbf{A}_1 \circ \mathbf{A}_2)$ that takes a basis \mathbf{S} of $\Lambda_q^\perp(\mathbf{A}_1)$ and a matrix \mathbf{A} with $\mathbb{Z}_q^{n \times m_1} \ni \mathbf{A}_1 \sqsubset \mathbf{A} \in \mathbb{Z}_q^{n \times (m_1+m_2)}$ as input. If $m_1 \geq 2n \log_2(q)$, it outputs a basis \mathbf{S} for $\Lambda_q^\perp(\mathbf{A})$ with $\|\tilde{\mathbf{S}}\| = \|\tilde{\mathbf{S}}_1\|$.*

Whenever trapdoor delegation is required, one cannot simply use extending control and hand over the resulting basis as it leaks information about the original trapdoor. Here, we can use tree propagation to obtain a *randomized* offspring with a new, random trapdoor.

Proposition 3 (Randomizing Control). *On input a basis \mathbf{S} of the lattice $\Lambda_q^\perp(\mathbf{A})$ of dimension m and a Gaussian parameter $s \geq \|\tilde{\mathbf{S}}\| \omega(\sqrt{\log(n)})$, the polynomial time algorithm $\text{RandBasis}(\mathbf{S}, s)$ outputs a basis \mathbf{S}' of $\Lambda_q^\perp(\mathbf{A})$ with $\|\tilde{\mathbf{S}}'\| \leq s\sqrt{m}$. The basis is independent of \mathbf{S} in the sense that for any two bases $\mathbf{S}_0, \mathbf{S}_1$ of $\Lambda_q^\perp(\mathbf{A})$ and $s \geq \max\{\|\tilde{\mathbf{S}}_0\|, \|\tilde{\mathbf{S}}_1\|\} \omega(\sqrt{\log(n)})$, $\text{RandBasis}(\mathbf{S}_0, s)$ is within negligible statistical distance of $\text{RandBasis}(\mathbf{S}_1, s)$.*

Estimating Secure Parameters. We could use the worst-case to average-case reduction for selecting secure parameters but that might be too conservative as the reduction is quite loose with respect to the ratio $m/n = \Omega(\log(n))$. The observations of Gama and Nguyen in [15] may be more realistic. They assume that lattice reduction algorithms find short lattice vectors \mathbf{v} in lattices Λ of dimension d with $\|\mathbf{v}\|_2 \leq \delta^d \det(\Lambda)^{1/d}$ for some $\delta > 0$. This value δ is supposed to “summarize” the capability of the employed lattice reduction algorithm. Nguyen and Gama analyze random lattices from a certain distribution and find that reaching $\delta < 1.01$ in high dimensions is hard.

Since the distribution of lattices in the SIS problem is different, their conjecture is not directly applicable. Furthermore, Micciancio and Regev describe an attack that works best in a sub-lattice of the usual q -ary lattice $\Lambda_q^\perp(\mathbf{A})$. Assuming an adversary with capability δ^* , they show that for given parameters n and q , the best strategy of for the adversary is to attack a sub-lattice of dimension $d^* = \sqrt{n \log(q) / \log(\delta^*)}$. In this dimension, the adversary finds vectors of Euclidean length $\nu^* = \min\{q, 2^2 \sqrt{n \log(q) \log(\delta^*)}\}$ [31]. For our parameter choices, we will always assume that the adversary is capable of reaching $\delta^* = 1.01$, but no less. All schemes in Sections 3 and 4 are based on the hardness of SIS in q -ary lattices of dimension m for a particular norm bound ν . Given this norm bound and δ^* , we need to establish that $\nu^* \gg \nu$, e.g., with a factor of 10 between both sides. The implicit assumption is that lattice basis reduction algorithms behave the same on random q -ary sub-lattices. Whether this heuristic is completely sound, is still unknown.

3 Warm-up — Constructions with Random Oracles

In this section, we recall strongly unforgeable GPV signatures as introduced in [16]. Then, we show how to use GPV together with the Bonsai-tree concept to build strongly unforgeable hierarchical identity-based signatures in the random oracle model. The proposed scheme is also secure under adaptive-identity queries.

3.1 Strongly Unforgeable Signatures

Lattice-based strongly unforgeable signatures were first proposed by Gentry, Peikert, and Vaikuntanathan in [16]. They introduce a family of preimage sampleable functions $\text{GPV} = (\text{TrapGen}, \text{Eval}, \text{SamplePre})$ on lattices. Its parameters $q, m, \tilde{L}, s = \omega(\sqrt{\log(n)})\tilde{L}$ only depend on the security parameter n as in Proposition 1. We define the sets $D_d := \{\mathbf{x} \in \mathbb{Z}^m \setminus \{\mathbf{0}\} : \|\mathbf{x}\|_2 \leq d\}$ and $R := \mathbb{Z}_q^n$.

The algorithm $\text{TrapGen}(1^n)$ outputs a public description $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a secret trapdoor $\mathbf{T} \in \mathbb{Z}^{m \times m}$, $\|\tilde{\mathbf{T}}\| \leq \tilde{L}$. Evaluation of the function $f_{\mathbf{A}} : \mathbb{Z}_q^m \rightarrow R$ is performed by $\text{Eval}(\mathbf{A}, \mathbf{x}) = \mathbf{A}\mathbf{x} \bmod q$. Finally, the inversion algorithm $\text{SamplePre}(\mathbf{T}, s, \mathbf{y})$ samples from the set of preimages $\{\mathbf{x} \in D_{s\sqrt{m}} : f_{\mathbf{A}}(\mathbf{x}) = \mathbf{y}\}$. Preimages have a conditional min-entropy of $\omega(\log(n))$ and follow a certain Gaussian distribution that can be efficiently sampled with SampleDom , even without the trapdoor. By construction, the function compresses the input and therefore admits collisions. Finding collisions in D_d , however, is at least as hard as solving $\text{SIS}(n, q, m, 2d)$.

Now, we can define DSig^{GPV} accordingly. Let \mathcal{H} be a family of random oracles $\mathbf{H} : \{0, 1\}^* \rightarrow R = \mathbb{Z}_q^n$. The key generator simply forwards the output of TrapGen . Signing a message \mathbf{m} involves choosing $r \leftarrow_{\mathfrak{s}} \{0, 1\}^n$, computing $\mathbf{y} \leftarrow \mathbf{H}(\mathbf{m}, r)$, and calling $\mathfrak{s} \leftarrow \text{SamplePre}(\mathbf{T}, s, \mathbf{y})$. A signature (\mathfrak{s}, r) for \mathbf{m} verifies if $\mathfrak{s} \in D_{s\sqrt{m}}$ and $f_{\mathbf{A}}(\mathfrak{s}) = \mathbf{H}(\mathbf{m}, r)$.

Security. Notice that the signature oracle can be efficiently simulated without the trapdoor by standard random oracle techniques. A forgery $(\mathbf{m}^*, \mathfrak{s}^*, r^*)$ in the SU-CMA experiment implies a collision $(\mathfrak{s}^*, \mathfrak{s}_i)$ with $f_{\mathbf{A}}(\mathfrak{s}_i) = f_{\mathbf{A}}(\mathfrak{s}^*) = \mathbf{H}(\mathbf{m}^*, r^*)$, where \mathfrak{s}_i was chosen during the random oracle simulation. By the conditional min-entropy, we know that $\mathfrak{s}_i \neq \mathfrak{s}^*$ with probability $1 - n^{-\omega(1)}$. Let $T_{\text{TrapGen}}(n)$, $T_{\text{SampleDom}}(n)$, and $T_{\text{Eval}}(n)$ be the cost functions for trapdoor generation, sampling, and trapdoor evaluation. In addition, let $T_{\text{List}}(n)$ be the cost function for list processing in the simulation of the random oracle.

Theorem 2 (Strong Unforgeability [16]). *For the above choice of parameters, DSig^{GPV} is (t, q_S, q_H, ϵ) -strongly unforgeable in the random oracle model if $\text{SIS}(n, m, q, 2s\sqrt{m})$ is $(t + q_S T_{\text{List}} + q_H(T_{\text{SampleDom}}(n) + T_{\text{Eval}}(n) + T_{\text{List}}(n)), (\epsilon - q_S^2/2^n)(1 - n^{-\omega(1)})$ -hard.*

Via Theorem 1, a corollary states that a successful attacker can find the shortest vector in all lattices of dimension n up to an approximation factor $\gamma \geq \tilde{L}\tilde{O}(n) = \tilde{O}(n\sqrt{n})$.

Secure Parameters. We estimate secure parameters for DSig^{GPV} as described in Section 2. For a given security parameter n , we choose the remaining parameters according to Proposition 1 and a prime $q \approx n^{4.5}$. For $n \geq 197$, we obtain the required hardness of the underlying SIS problem.

3.2 Strongly Unforgeable Hierarchical ID-based Signatures

Using the signature scheme from the previous section, we can directly apply the Bonsai-tree concept to obtain a hierarchical identity-based signature scheme HIBS^{GPV} in the random oracle model. The idea is that, based on a given ID, the secret key extraction algorithm `Extract` first uses `ExtBasis` to extend the matrix \mathbf{A}^* to \mathbf{A}_{ID} and the secret master key \mathbf{T}^* to \mathbf{T}_{ID} such that $\mathbf{A}_{\text{ID}}\mathbf{T}_{\text{ID}} \equiv \mathbf{0}$. Then, in order to protect the master key, it uses `RandBasis` and outputs a randomized trapdoor \mathbf{S}_{ID} . The individual signers use DSig^{GPV} and we have an identity-based signature scheme that is strongly unforgeable. This concept can be transferred to the hierarchical setting, where every signer may act as a key extraction entity. Note that the number of levels in the hierarchy affects the tightness of the security proof as these randomized trapdoors are slightly worse than the master trapdoor.

We assume that all identities on all levels have length κ . The maximum number of levels, including the master-key, is $\ell + 1$. We denote the basis length on level k with \tilde{L}_k and the corresponding Gaussian parameter with $s_k = \omega(\sqrt{\log(n)})\tilde{L}_k$.

Master-key Generation. Let q, \tilde{L}, m_1, m_2 be chosen according to Proposition 1 and let $d = \tilde{L}\omega(\sqrt{\log(n)})^{\ell+1}\sqrt{m_1 + (\ell\kappa + 1)m_2}^{-\ell+1}$. These parameters may be excluded from the public key as they are the same for all users. Generate a description $\mathbf{A}^* \in \mathbb{Z}_q^{n \times m_1 + m_2}$ of the master lattice $\Lambda_q^\perp(\mathbf{A}^*)$ together with a trapdoor \mathbf{S}^* such that $\|\tilde{\mathbf{S}}\| \leq \tilde{L}$ with `ExtLattice`. Then, choose random

matrices $\langle \mathbf{A} \rangle := \left\{ (\mathbf{A}_i^{(0)}, \mathbf{A}_i^{(1)}) \right\}_1^{\ell\kappa}$ from $\mathbb{Z}_q^{n \times m_2}$. The secret is \mathbf{S}^* and the public key is $(\mathbf{A}^*, \langle \mathbf{A} \rangle)$.

Key Extraction. On input $\mathbf{S}_{\text{ID}^*}^*$, ID with $\text{ID} = \text{ID}^* \circ \text{ID}'$, $|\text{ID}| = l \leq \ell\kappa$, $\text{ID}' = \text{ID}'_1, \dots, \text{ID}'_\kappa \in \{0, 1\}^\kappa$ recursively define the matrix $\mathbf{A}_{\text{ID}} := \mathbf{A}_{\text{ID}^*} \circ \mathbf{A}_{\text{ID}'_1}^{(\text{ID}'_1)} \circ \dots \circ \mathbf{A}_{\text{ID}'_\kappa}^{(\text{ID}'_\kappa)}$ with $\mathbf{A}_\emptyset := \mathbf{A}^*$, and call $\mathbf{T}_{\text{ID}} \leftarrow \text{ExtBasis}(\mathbf{S}_{\text{ID}^*}^*, \mathbf{A}_{\text{ID}})$. Then, let $s = \|\tilde{\mathbf{T}}_{\text{ID}}\| \omega(\sqrt{\log(n)})$ and output the randomized basis $\mathbf{S}_{\text{ID}} \leftarrow \text{RandBasis}(\mathbf{S}_{\text{ID}}, s)$ with $\|\tilde{\mathbf{S}}_{\text{ID}}\| \leq s\sqrt{\dim(\mathbf{S}_{\text{ID}})}$. For inappropriate inputs, return \perp .

Signing. On input a message $\mathbf{m} \in \{0, 1\}^*$ and the trapdoor \mathbf{S}_{ID} , the signer with identity ID on level k chooses $r \leftarrow_{\mathfrak{s}} \{0, 1\}^n$ and computes $\mathfrak{s} \leftarrow \text{SamplePre}(\mathbf{S}_{\text{ID}}, s_k, \mathbf{H}(\mathbf{m}, r, \text{ID}))$. It outputs (\mathfrak{s}, r) .

Verification. On input $(\mathbf{A}^*, \langle \mathbf{A} \rangle)$, an identity ID , a signature (\mathfrak{s}, r) , and a message \mathbf{m} , the algorithm outputs 1 if and only if $\mathfrak{s} \in D_d$ and $f_{\mathbf{A}_{\text{ID}}}(\mathfrak{s}) = \mathbf{H}(\mathbf{m}, r, \text{ID})$.

Recall that SamplePre outputs a vector of length at most $s\sqrt{m}$ with $s = \omega(\sqrt{\log(n)})$ \tilde{L} when called with a trapdoor of length at most \tilde{L} in dimension m . The maximum dimension in the scheme is $m_1 + (\ell\kappa + 1)m_2$ on level ℓ . The maximum basis length on level i is $\tilde{L}_i = s_{i-1}\sqrt{m_1 + (\ell\kappa + 1)m_2}$ with $\tilde{L}_0 := \tilde{L}$. The Gaussian parameter for presampling on level i is $s_i = \omega(\sqrt{\log(n)})\tilde{L}_i$ with $s_0 = s$. This recurrence yields $\tilde{L}_\ell = \tilde{L}\omega(\sqrt{\log(n)})^\ell \sqrt{m_1 + (\ell\kappa + 1)m_2}^\ell$. Signing with such a basis yields signatures of Euclidean length at most $\omega(\sqrt{\log(n)})\tilde{L}_\ell \sqrt{m_1 + (\ell\kappa + 1)m_2} = d$. Thus, such signatures are accepted by the verifier. The security guarantees scale with the depth of the hierarchy as the norm bound becomes looser for increasing ℓ .

Security. We prove that HIBS^{GPV} is secure under selective identity attacks. The second theorem shows that it is even secure under *adaptive* identity attacks, but with a looser reduction. The latter reduction, however, is not as loose as the generic method of guessing the right identity with probability $2^{-\kappa}$ [8]. Let $T_{\text{func}}(x)$ be the cost function for function func and let $T_{\text{list}}(n)$ be the cost function for list processing for simulating a random oracle.

Theorem 3 (Selective Security). HIBS^{GPV} is (t, q_S, q_H, ϵ) -strongly unforgeable under selective identity attacks in the random oracle model if SIS is $(t + 2T_{\text{ExtLattice}} + q_E T_{\text{Extract}} + (q_H + q_S)T_H, (1 - n^{-\omega(1)})(\epsilon - q_S^2/2^n))$ -hard with norm bound $\nu = 2\tilde{L}\omega(\sqrt{\log(n)})^{\ell+1} \sqrt{m_1 + (\ell\kappa + 1)m_2}^{\ell+1}$.

Notice the the GPV signature scheme can be efficiently simulated by a standard random oracle technique. Moreover, the adversary will make $\leq q_E$ extraction queries that need to be answered. We prepare for this by “knowing” a trapdoor for a prefix of all but the challenge identity. Upon an extraction query, this trapdoor can be extended to a trapdoor for the requested identity. The external challenge, the input \mathbf{A} from the SIS problem, is embedded in the public key of the challenge identity. Via random oracle techniques, the reduction knows a

valid signature for the output message of the adversary. However, the adversary outputs a different signature with probability $1 - n^{-\omega(1)}$ by the conditional min-entropy of the set of possible signature. The proof is in the full version [33].

Theorem 4 (Adaptive Security). *HIBS^{GPV} can be made $(t, q_S, q_H, q_G, \epsilon)$ -strongly unforgeable under adaptive identity attacks in the random oracle model if it is $(t, q_S, q_H, q_G, \epsilon/q_G)$ -strongly unforgeable under selective identity attacks.*

Here, we only give the idea of the conversion. We need to change the way the identities are mapped to the branches of the Bonsai tree. Instead of directly using the binary representation of ID or its hash value, we apply a random oracle G to the individual substrings of length κ first. Thus, every ID is mapped to a random position in the tree. Assume that the adversary makes q_G queries to this random oracle, we can prepare a randomly selected path in the tree with “undirected growth” and program the random oracle to map one of the q_G queries to this path. Therefore, the success probability of the reduction degrades with a factor $1/q_G$ instead of the generic $1/2^\kappa$.

The worst-case to average-case reduction guarantees security if finding shortest vectors up to an approximation factor $\gamma \geq 2d\tilde{O}(\sqrt{n}) = \tilde{O}(\sqrt{n}^{\ell+3})$ is hard in the worst case in dimension n .

Secure Parameters. We assume identities of length $\kappa = 40$ bits, which should be sufficient in practical scenarios. The parameter q is crucial as it greatly influences the hardness of SIS and the worst-case to average-case reduction only holds for large enough q . The scheme’s efficiency greatly depends on the number $\ell + 1$ of layers in the hierarchy. For larger ℓ , we are forced to increase q . Possible secure choices for (ℓ, n, q) are $(1, 467, n^6)$, $(2, 692, n^8)$, $(3, 1011, n^9)$, or $(4, 1258, n^{11})$.

4 Constructions without Random Oracles

In the following, we propose our main constructions. We start with a strongly unforgeable signature scheme and then propose a strongly unforgeable hierarchical identity-based signature scheme that is secure under selective identity attacks. Both constructions are secure in the standard model.

In the following, we let $H : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ be a collision-resistant hash function that maps into the message space. Finding collisions in polynomial time is only possible with probability $\leq c$.

4.1 Strongly Unforgeable Signatures

We propose a variant of the EU-CMA signature scheme [12]. The authors first create an EU-SMA secure signature scheme in the standard model and then apply a generic conversion from EU-SMA to EU-CMA, using a Chameleon hash function. We follow this line of thought and construct a signature scheme that is even secure in the SU-CMA sense. We achieve this by computing the signatures

in a different way, compared to [12] scheme. In their scheme, there is a matrix \mathbf{A}_m and a corresponding lattice $\Lambda_q^\perp(\mathbf{A}_m)$ for which the signer can derive a trapdoor. The matrix \mathbf{A}_m is formed as in Section 3.2. The signatures in their work are short lattice vectors in $\Lambda_q^\perp(\mathbf{A}_m)$, i.e., short vectors $\mathfrak{s} \neq \mathbf{0}$ such that $\mathbf{A}_m \mathfrak{s} \equiv \mathbf{0}$. In our scheme, we fix a random $\mathbf{y} \in \mathbb{Z}_q^n$ and let the signer solve $\mathbf{A}_m \mathfrak{s} \equiv \mathbf{y}$ instead. The overhead is a simple linear algebra step (cf. [16]) that can be done once during key generation. Surprisingly, this slight change enables us to prove *strong* unforgeability. It is important to note that, unlike in [12], we need to be able to answer all signature queries in the security proof, i.e., even for the message \mathbf{m}^* that the adversary outputs a forgery for.

First of all, we demonstrate that the scheme in [12] is not strongly unforgeable by showing an attack that only works in the SU-SMA model and not in EU-SMA. The adversary in the SU-SMA experiment queries its signature oracle with a random message \mathbf{m}^* and receives a signature \mathfrak{s} such that $\mathfrak{s} \in D(s\sqrt{m})$ and $\mathbf{A}_m \mathfrak{s} \equiv \mathbf{0}$. The adversary simply returns the valid forgery $\mathfrak{s}^* \leftarrow -\mathfrak{s}$.

We specify $\text{DSig}^{\text{su-sma}}$, where one may interpret the randomness-message pair as an identity.

Key Generation. Let q, \tilde{L}, m_1, m_2 be chosen according to Proposition 1 and let $s = \tilde{L}\omega(\sqrt{\log(n)})$ and $d = s\sqrt{m_1 + (\lambda + 1)m_2}$. These parameters may be excluded from the public key as they are the same for all users. Use ExtLattice to generate a description $\mathbf{A}^* \in \mathbb{Z}_q^{n \times m_1 + m_2}$ of the master lattice

$\Lambda_q^\perp(\mathbf{A}^*)$ together with a trapdoor \mathbf{S}^* such that $\|\tilde{\mathbf{S}}\| \leq \tilde{L}$. Furthermore,

generate a set $\langle \mathbf{B} \rangle := \left\{ (\mathbf{B}_i^{(0)}, \mathbf{B}_i^{(1)}) \right\}_1^\lambda$ random matrices in $\mathbb{Z}_q^{n \times m_2}$. Finally, choose $\mathbf{y} \leftarrow_{\mathfrak{s}} \mathbb{Z}_q^n$ and output the secret key \mathbf{S}^* and the public key $(\mathbf{A}^*, \langle \mathbf{B} \rangle, \mathbf{y})$.

Signing. On input a message $\mathbf{m} \in \{0, 1\}^*$ and the secret trapdoor \mathbf{S}^* , the signer with identity chooses $r \leftarrow_{\mathfrak{s}} \{0, 1\}^n$ and computes $h \leftarrow \mathbf{H}(\mathbf{m}, r)$, $\mathfrak{s} \leftarrow \text{SamplePre}(\mathbf{S}_h, s, \mathbf{y})$. The trapdoor \mathbf{S}_h is formed via $\text{ExtBasis}(\mathbf{S}^*, \mathbf{A}_h)$, where $\mathbf{A}_h := \mathbf{A}^* \circ \mathbf{B}_1^{(h_1)} \circ \dots \circ \mathbf{B}_\lambda^{(h_\lambda)}$. It outputs (\mathfrak{s}, r) .

Verification. On input $(\mathbf{A}^*, \langle \mathbf{B} \rangle, \mathbf{y})$, a signature (\mathfrak{s}, r) , and a message \mathbf{m} , the algorithm outputs 1 if and only if $\mathfrak{s} \in D_d$ and $\mathbf{A}_{\mathbf{H}(\mathbf{m}, r)}(\mathfrak{s}) = \mathbf{y}$.

The scheme is complete because all signatures are generated using a basis of length \tilde{L} and with the Gaussian parameter $s = \omega(\sqrt{\log(n)})\tilde{L}$. The total dimension is $m = m_1 + (\lambda + 1)m_2$. Thus, SamplePre outputs signatures of length at most $s\sqrt{m_1 + (\lambda + 1)m_2} = d$ that are accepted by Vf . In order to get the full (SU-CMA) scheme, we wrap the message with a Chameleon hash function.

Security. We prove that $\text{DSig}^{\text{su-sma}}$ is SU-SMA secure and then apply the black-box conversion in Lemma 1. Let $T_{\text{func}}(n)$ be the cost function for the function func and let $T_{\text{list}}(n)$ be the cost function for list processing, which is explained in the analysis below.

Theorem 5 (Strong Unforgeability). $\text{DSig}^{\text{su-sma}}$ is (t, ϵ) strongly unforgeable under static message attacks (SU-SMA) if SIS with $\nu = 2\tilde{L}\omega(\sqrt{\log(n)})$

$\sqrt{m_1 + (\lambda + 1)m_2}$ is $(t + \lambda q_S T_{\text{List}} + T_{\text{ExtLattice}} + q_S(T_{\text{SamplePre}} + T_{\text{ExtBasis}}), 1/2(1 - n^{-\omega(1)})(\epsilon - q_S^2/2^n - c)/(\lambda q_S))$ -hard.

The idea is to separate the adversaries into two classes. One works in the EU-SMA sense and the other exploits the additional freedom of the SU-SMA setting. The reduction guesses the type of adversary before handing over the public key. If it expects an EU-SMA forger, the reduction knows \mathbf{x} with $\mathbf{A}^* \mathbf{x} \equiv \mathbf{y}$ and forces the forger to solve an inhomogeneous SIS, for which the reduction does not know the trapdoor. Together with \mathbf{x} , it can solve the corresponding SIS with overwhelming probability. For the SU-SMA forger, the reduction has to guess the index i^* of the signature query that will be recycled in the forgery with probability $1/q_S$. There, it plants an \mathbf{x} with $\mathbf{A}_{\mathbb{H}(\mathbf{m}_{i^*}, r_{i^*})} \mathbf{x} \equiv \mathbf{y}$. Again, with the adversary's help, the reduction solves SIS with overwhelming probability, while being able to answer a single signature query for \mathbf{m}_{i^*} with \mathbf{x} . The key extraction queries are answered as described in Section 3.2. The full proof is in the full version [33]

Therefore, $\text{DSig}^{\text{su-sma}}$ is secure if finding shortest vectors in dimension n , within factors $\gamma \geq \tilde{O}(n\sqrt{n})$, is hard in the worst case. Via Lemma 1, we immediately get similar results for SU-CMA.

Secure Parameters. The underlying problem is SIS with $\nu = 2\omega(\sqrt{\log(n)})\tilde{L}\sqrt{m_1 + (\lambda + 1)m_2}$. Let $\lambda = 160$ and consider $\log(n) = \omega(\sqrt{\log(n)})$. As before, q is chosen such that the worst-case to average-case reduction holds ($q \approx n^5$). Now, for $n \geq 247$, we obtain the required complexity of SIS.

4.2 Strongly Unforgeable Hierarchical ID-based Signatures

By adding more layers to the hierarchy in Section 4.1, we construct a hierarchical identity-based signature scheme $\text{HIBS}^{\text{su-sma}}$ that is strongly unforgeable. The identities are handled as in Section 3.2 but then we use an additional Bonsai tree to sign.

The length of each sub-identity on each level is κ , the number of levels is ℓ , and messages have λ bits. \tilde{L}_k is the basis length on level k . The corresponding Gaussian parameter is $s_k := \omega(\sqrt{\log(n)})\tilde{L}_k$.

Master-key Generation. Let q, \tilde{L}, m_1, m_2 be chosen according to Proposition 1 and let $d = \tilde{L}\omega(\sqrt{\log(n)})^{\ell+1}\sqrt{m_1 + (\ell\kappa + \lambda + 1)m_2}^{\ell+1}$. These parameters may be excluded from the public key as they are the same for all users. Use ExtLattice to generate a description $\mathbf{A}^* \in \mathbb{Z}_q^{n \times m_1 + m_2}$ of the master lattice $\Lambda_q^\perp(\mathbf{A}^*)$ together with a trapdoor \mathbf{S}^* such that $\|\tilde{\mathbf{T}}\| \leq \tilde{L}$. Furthermore, generate the sets $\langle \mathbf{A} \rangle := \left\{ (\mathbf{A}_i^{(0)}, \mathbf{A}_i^{(1)}) \right\}_1^{\ell\kappa}$, $\langle \mathbf{B} \rangle := \left\{ (\mathbf{B}_i^{(0)}, \mathbf{B}_i^{(1)}) \right\}_1^\lambda$ of random matrices in $\mathbb{Z}_q^{n \times m_2}$. Finally, choose $\mathbf{y} \leftarrow_{\S} \mathbb{Z}_q^n$ and output the secret key \mathbf{S}^* and the public key $(\mathbf{A}^*, \langle \mathbf{A} \rangle, \langle \mathbf{B} \rangle, \mathbf{y})$.

Key Extraction. On input $\mathbf{S}_{\text{ID}^*}^*$, ID with $\text{ID} = \text{ID}^* \circ \text{ID}'$, $|\text{ID}| = l \leq \ell\kappa$, $\text{ID}' = \text{ID}'_1, \dots, \text{ID}'_\kappa \in \{0, 1\}^\kappa$ recursively define the matrix $\mathbf{A}_{\text{ID}} := \mathbf{A}_{\text{ID}^*} \circ \mathbf{A}_{l+1}^{(\text{ID}'_1)} \circ \dots \circ \mathbf{A}_{l+\kappa}^{(\text{ID}'_\kappa)}$ with $\mathbf{A}_\emptyset := \mathbf{A}^*$, and call $\mathbf{T}_{\text{ID}} \leftarrow \text{ExtBasis}(\mathbf{S}_{\text{ID}^*}^*, \mathbf{A}_{\text{ID}})$. Then, let $s = \|\tilde{\mathbf{T}}\| \omega(\sqrt{\log(n)})$ and output the randomized basis $\mathbf{S}_{\text{ID}} \leftarrow \text{RandBasis}(\mathbf{T}_{\text{ID}}, s)$ with $\|\tilde{\mathbf{S}}_{\text{ID}}\| \leq s\sqrt{\dim(\mathbf{T}_{\text{ID}})}$. For inappropriate inputs, return \perp .

Signing. On input a message $\mathbf{m} \in \{0, 1\}^*$ and a secret trapdoor \mathbf{S}_{ID} , the signer with identity ID on level k chooses $r \leftarrow_{\mathfrak{s}} \{0, 1\}^n$ and computes $\mu \leftarrow \text{H}(\mathbf{m}, r)$. Now μ is used to form $\mathbf{A}_{\text{ID} \circ \mu} := \mathbf{A}_{\text{ID}} \circ \mathbf{B}_1^{(\mu_1)} \circ \dots \circ \mathbf{B}_\lambda^{(\mu_\lambda)}$. Then, the signer extends its secret basis by calling $\mathbf{S}_{\text{ID} \circ \mu} \leftarrow \text{ExtBasis}(\mathbf{S}_{\text{ID}}, \mathbf{A}_{\text{ID} \circ \mu})$. Finally, it outputs $\mathfrak{s} \leftarrow \text{SamplePre}(\mathbf{S}_{\text{ID} \circ \mu}, s_k, \mathbf{y})$ and r .

Verification. On input $(\mathbf{A}^*, \langle \mathbf{A} \rangle, \langle \mathbf{B} \rangle, \mathbf{y})$, a signature (\mathfrak{s}, r) , and a message \mathbf{m} , the algorithm outputs 1 if and only if $\mathfrak{s} \in D_d$ and $\mathbf{A}_{\text{ID} \circ \text{H}(\mathbf{m}, r)}(\mathfrak{s}) = \mathbf{y}$.

Notice that the scheme is complete by a similar argument as in Section 3.2. The maximum trapdoor length (level ℓ) is $\tilde{L}_\ell = \tilde{L}\omega(\sqrt{\log(n)})^\ell \sqrt{m_1 + (\ell\kappa + \lambda + 1)m_2}$. This basis is extended once before signing, which is length-preserving. Then, signing is done via SamplePre that outputs a vector of length at most $\omega(\sqrt{\log(n)}) \tilde{L}_\ell \sqrt{m_1 + (\ell\kappa + \lambda + 1)m_2} \leq d$. Thus, the verification algorithm accepts.

Unforgeability. We prove that $\text{HIBS}^{\text{su-sma}}$ is SU-SMA secure under a selective-identity attack. Then, we can apply the transform in Lemma 1 to make it SU-CMA secure. For the transformation to be identity-based as well, the Chameleon hash function needs to be published as part of mpk .

Theorem 6 (Selective Security). $\text{HIBS}^{\text{su-sma}}$ is (t, q_S, ϵ) SU-SMA secure under selective identity attacks if SIS is $(t + 3T_{\text{ExtLattice}} + q_E T_{\text{Extract}} + \lambda q_S T_{\text{List}} + q_S(T_{\text{SamplePre}} + T_{\text{ExtBasis}}), 1/2(1 - n^{-\omega(1)})(\epsilon - q_S^2/2^n - c)/(\lambda q_S))$ -hard with norm bound $\nu = 2\tilde{L}\omega(\sqrt{\log(n)})^{\ell+1} \sqrt{m_1 + (\ell\kappa + \lambda + 1)m_2}$.

Proof (Sketch). We assume that there is a successful adversary \mathcal{A} against unforgeability of $\text{HIBS}^{\text{su-sma}}$ and construct a reduction \mathcal{B} that solves SIS. The setup is as in Theorem 3 in order to be able to simulate the key extraction queries for all identities that are not a prefix of the challenge identity I . However, the input \mathbf{A} of the reduction needs to be wider, namely in $\mathbb{Z}_q^{m_1 + (\ell\kappa + \lambda + 1)m_2}$ in order to simulate the static signature queries for identity I . Signature queries for identities $I' \neq I$ can be simulated with a known trapdoor. The overhead of the reduction is $2T_{\text{ExtLattice}} + q_E T_{\text{Extract}}$ for setting up and running the extraction oracle. The overhead for signing is $\lambda q_S T_{\text{List}} + T_{\text{ExtLattice}} + q_S(T_{\text{SamplePre}} + T_{\text{ExtBasis}})$. As in Theorem 5, the reduction sets \mathbf{y} such that it knows a preimage \mathbf{x} either for EU-SMA or for SU-SMA forgers. In both cases, the reduction solves SIS. The probability that the adversary outputs a signature that yields a solution to SIS is as in Theorem 5 because we use the same signature scheme here and we can prepare for all extraction queries.

The full proof is a combination of the techniques that are used in Theorems 3 and 5. \square

In consequence, $\text{HIBS}^{\text{su-sma}}$ is secure as long as finding shortest lattice vectors up to approximation factors $\gamma \geq 2d\tilde{\mathcal{O}}(\sqrt{n}) = \tilde{\mathcal{O}}(n\sqrt{n}^{\ell+3})$ is hard in the worst case in dimension n .

Secure Parameters. We let $\kappa = 40$ and $\lambda = 160$ and estimate secure parameters for $\ell = 1, 2, 3, 4$, i.e., for 2, 3, 4, 5 levels. For (ℓ, n, q) , we obtain $(1, 478, n^7)$, $(2, 730, n^9)$, $(3, 1082, n^{10})$, or $(4, 1362, n^{12})$.

5 Conclusions

We have shown three results. As a warm-up, we have constructed an adaptive-ID secure, strongly unforgeable hierarchical identity-based signature scheme in the random oracle model. Then, we have shown the first lattice-based strongly unforgeable signature scheme without Merkle trees [29] in the standard model. And, finally, we provide the first strongly unforgeable hierarchical identity-based signatures scheme in the standard model from lattices. For each construction, we have proposed a set of secure parameters based on today's knowledge about lattice reduction algorithms. Another benefit is that all of our constructions can be transferred to ideal lattices that admit short keys and efficient operations. Moreover, we can use mild, worst-case hardness assumptions in lattices as the basis of security in all constructions.

Acknowledgments

The author would like to thank Benoît Libert for a helpful discussion on HIBS. He also thanks Pierre-Louis Cayrel and Dominique Schröder for reviewing parts of this work. Furthermore, the author thanks the anonymous reviewers of PQC 2010 for their constructive comments. The author is indebted to one of the reviewers, who suggested a simplification for demonstrating that the scheme in [12] is not SU-CMA secure. This part of Section 4 is now much easier to understand.

References

1. Shweta Agrawal and Xavier Boyen. Identity-based encryption from lattices in the standard model. Manuscript, July 2009. Available at <http://www.cs.stanford.edu/~xb/ab09/>.
2. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108. ACM, 1996.
3. Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610. ACM, 2001.
4. Joel Alwen and Chris Peikert. Generating shorter bases for hard random lattices. Cryptology ePrint Archive, Report 2008/521, 2008. <http://eprint.iacr.org/>.

5. Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. In Susanne Albers and Jean-Yves Marion, editors, *STACS*, volume 09001 of *Dagstuhl Seminar Proceedings*, pages 75–86. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2009.
6. Mihir Bellare and Sarah Shoup. Two-tier signatures, strongly unforgeable signatures, and fiat-shamir without random oracles. In Tatsuaki Okamoto and Xiaoyun Wang, editors, *Public Key Cryptography*, volume 4450 of *Lecture Notes in Computer Science*, pages 201–216. Springer, 2007.
7. Daniel J. Bernstein, Johannes A. Buchmann, and Erik Dahmen, editors. *Post-Quantum Cryptography*. Springer, 2008.
8. Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In Christian Cachin and Jan Camenisch, editors, *Eurocrypt*, volume 3027 of *LNCS*, pages 223–238. Springer, 2004.
9. Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput.*, 36(5):1301–1328, 2007.
10. Dan Boneh, Emily Shen, and Brent Waters. Strongly unforgeable signatures based on computational diffie-hellman. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 229–240. Springer, 2006.
11. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, 2004.
12. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. to appear in EUROCRYPT 2010.
13. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM J. Comput.*, 30(2):391–437, 2000.
14. David Galindo, Javier Herranz, and Eike Kiltz. On the generic construction of identity-based signatures with additional properties. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT*, volume 4284 of *Lecture Notes in Computer Science*, pages 178–193. Springer, 2006.
15. Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In Nigel P. Smart, editor, *EUROCRYPT*, volume 4965 of *LNCS*, pages 31–51. Springer, 2008.
16. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *STOC*, pages 197–206. ACM, 2008.
17. Craig Gentry and Alice Silverberg. Hierarchical id-based cryptography. In Yuliang Zheng, editor, *ASIACRYPT*, volume 2501 of *Lecture Notes in Computer Science*, pages 548–566. Springer, 2002.
18. Shai Halevi, editor. *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *LNCS*. Springer, 2009.
19. Susan Hohenberger and Brent Waters. Short and stateless signatures from the rsa assumption. In Halevi [18], pages 654–670.
20. Eike Kiltz, Anton Mityagin, Saurabh Panjwani, and Barath Raghavan. Append-only signatures. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP*, volume 3580 of *LNCS*, pages 434–445. Springer, 2005.
21. Eike Kiltz and Gregory Neven. Identity-based signatures. In M. Joye and G. Neven, editors, *Cryptology and Information Security Series*, volume 2, pages 31–44. IOS Press, 2008.

22. Hugo Krawczyk and Tal Rabin. Chameleon hashing and signatures. Cryptology ePrint Archive, Report 1998/010, 1998. <http://eprint.iacr.org/>.
23. Hugo Krawczyk and Tal Rabin. Chameleon signatures. In *NDSS*. The Internet Society, 2000.
24. Gaëtan Leurent and Phong Q. Nguyen. How risky is the random-oracle model? In Halevi [18], pages 445–464.
25. Benoît Libert and Jean-Jacques Quisquater. The exact security of an identity based signature and its applications. Cryptology ePrint Archive, Report 2004/102, 2004. <http://eprint.iacr.org/>.
26. Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In Matsui [28], pages 598–616.
27. Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In Ran Canetti, editor, *TCC*, volume 4948 of *LNCS*, pages 37–54. Springer, 2008.
28. Mitsuru Matsui, editor. *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*. Springer, 2009.
29. Ralph C. Merkle. A certified digital signature. In Gilles Brassard, editor, *CRYPTO*, volume 435 of *LNCS*, pages 218–238. Springer, 1989.
30. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Prelim. in FOCSS 2002.
31. Daniele Micciancio and Oded Regev. Lattice-based cryptography. In Bernstein et al. [7], pages 147–191.
32. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.
33. Markus Rückert. Strongly unforgeable signatures and hierarchical identity-based signatures from lattices without random oracles. Cryptology ePrint Archive, Report 2010/070, 2010. <http://eprint.iacr.org/>.
34. Adi Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.
35. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Matsui [28], pages 617–635.